Lucas vs. Lucas: On Inequality and Growth*

Juan Carlos Córdoba†

and

Geneviève Verdier‡

November 17, 2005

Abstract

Lucas (2004) asserts that "Of the tendencies that are harmful to sound economics, the most seductive, and in my opinion the most poisonous, is to focus on questions of distribution...The potential for improving the lives of poor people by finding different ways of distributing current production is *nothing* [Italics in the original] compared to the apparently limitless potential of increasing production." In this article we evaluate this claim using an extended version of Lucas' (1987) welfare evaluation framework. We construct a social welfare function following Lucas' (2004) own suggestion of weighing everyone's welfare equally, and compute welfare measures in the same way as Lucas (1987). The result is surprising and robust. The potential welfare gains of redistribution are substantial and likely exceed the welfare gains of economic growth. Moreover, our calculations suggest that US inequality is above its optimal level.

Keywords: Welfare costs, business cycles, economic growth, inequality.

JEL Classification: E1, E2, D3

---

*We would like to thank Rui Castro, Dror Goldberg, Marvin Goodfriend, David Levine, Adi Mayer, Camille Nelson and seminar participants at the Universities of Pittsburgh, Texas A&M, and Rochester for helpful comments and suggestions. This research was conducted while the first author visited Carnegie Mellon University. Any errors or omissions are our own.

†Department of Economics, Rice University, 6100 Main Street, Houston TX, 77005, e-mail: `jcordoba@rice.edu`.

‡Department of Economics, Texas A&M University, 4228 TAMU, College Station TX, 77843-4228, e-mail: `gverdier@econmail.tamu.edu`.

1 Introduction

In the 2003 Annual Report of the Federal Reserve Bank of Minneapolis, Robert Lucas asserts:

> "Of the tendencies that are harmful to sound economics, the most seductive, and in my opinion the most poisonous, is to focus on questions of distribution. In this very minute, a child is being born to an American family and another child, *equally valued by God* [italics ours], is being born to a family in India. The resources of all kinds that will be at the disposal of this new American will be on the order of 15 times the resources available to his Indian brother. This seems to us a terrible wrong, justifying direct corrective action, and perhaps some actions of this kind can and should be taken. But of the vast increase in the well-being of hundreds of millions of people that has occurred in the 200-year course of the industrial revolution to date, virtually none of it can be attributed to the direct redistribution of resources from rich to poor. The potential for improving the lives of poor people by finding different ways of distributing current production is *nothing* [Italics in the original] compared to the apparently limitless potential of increasing production."
> Lucas (2004)

Although the positive statements in this assertion are likely uncontroversial, the normative conclusion is not. Lucas stresses the overwhelming importance of economic growth and dismisses redistributive issues as 'poisonous' and quantitatively unimportant, a conclusion that is based on the observation that economic growth helps to eliminate poverty. However, a correct assessment of the potential benefits of redistributive policies must be based on proper welfare analysis, such as the one used by Lucas (1987) to evaluate the costs of business cycles and benefits of economic growth. Welfare analysis takes into account features of individual preferences, in particular time discounting, aversion to consumption risk and dispersion, and inequality levels; factors that may significantly enhance the role of redistribution for social welfare and downplay the role of economic growth.

The view that economic growth is the major determinant of social welfare is partly supported by Lucas' own calculations. Lucas (1987) shocked the profession by showing how insignificant the potential welfare benefits of eliminating business cycles are, particularly when compared to the potential gains of increasing economic growth. Lucas' (2004) view on the costs of inequality and benefits of growth has profound implications. If right, it provides quantitative support to eliminate or limit costly redistributive programs and institutions. Moreover, if economic growth is all that matters for aggregate welfare, then societies should embrace growth-enhancing institutions and policies, ignoring potentially adverse but likely minor distributive consequences. One might then wonder: "Why do societies choose or end up with institutions that do not maximize economic growth or aggregate economic welfare?" (Acemoglu 2005, page 1)

Lucas' view however, is not shared by all. At least since Ricardo, distributional issues have been considered of the utmost importance in economics. This view, for example, was expressed by Ricardo in a famous letter to Malthus:

> "Political Economy you think is an enquiry into the nature and causes of wealth - I think it should rather be called an enquiry into the laws which determine the division of the produce industry amongst the classes who concur in its formation. No law can be laid down respecting quantity, but a tolerably correct one can be laid down respecting proportions. Every day I am more satisfied that the former enquiry is vain and delusive, and the latter only the true objects of the science." (Ricardo, 1951, pp. 278-279)[1]

This article evaluates Lucas' (2004) assertion using a version of Lucas' (1987) welfare evaluation framework. As mentioned above, the distinguishing feature of welfare analysis is that cost and gains are evaluated in present value terms, attitudes toward consumption risk and dispersion play a role, and the degree of inequality matters. We further assess optimal inequality and optimal growth by introducing technological restrictions — a production possibility frontier for inequality, consumption levels and growth — into the basic welfare framework. We find that the potential welfare gains of inequality reduction are far from *nothing*, and that they may actually exceed the potential gains of economic growth. Such an important role of inequality for welfare implies that maximizing social welfare is not equivalent to maximizing economic growth. In fact, our findings suggest that actual institutional choices reflect, to some important extent, a compromise between

[1]Ricardo's view is partly explained, as Lucas (2002) points out, by the fact classical economists did not suspect that economic growth would last.

inequality and efficiency. This, of course, is an old idea (see for example, Okun 1975) but it relies on the premise that inequality has a major effect on aggregate welfare, which is precisely what our evaluation reveals.

Lucas (1987) studies the welfare of a representative consumer characterized by an isoelastic utility function and a log-normal distribution of consumption. He defines welfare costs in terms of permanent compensation rates on consumption required to leave the consumer indifferent between the current situation and an alternative ideal situation.[2] Lucas reports measures about the cost of business cycles for different degrees of intertemporal substitutability, but unfortunately only reports welfare measures of economic growth for the log-utility case.

We extend this framework in three ways: we evaluate welfare gains of growth not only for the log case but also when the intertemporal elasticity of substitution differs from one; we introduce consumption inequality across individuals; and finally, we introduce technological restrictions.

First, we consider different degrees of intertemporal substitutability to assess the social benefits and costs of inequality and growth. This simple extension is important because the existing empirical literature typically regards the intertemporal elasticity of substitution (IES) to be lower than one, and welfare measures depend significantly on the size of this elasticity.[3] For example, the welfare gains of economic growth are lower than the costs of business cycles if the IES is sufficiently low. For commonly used values of the IES, however, the gains of economic growth are significant but perhaps smaller than previously thought. We find that the welfare gains of one additional point of economic growth per year range from 2% to 21% for the various values of the IES considered by Lucas (1987). These gains are typically lower than the 20% reported by Lucas. For a per-capita consumption growth rate of 2.1% — the 1960-2000 average for the 108 economies analyzed here — the welfare gains of total economic growth range from 7.6% to 51%.

The second extension is to introduce consumption dispersion across individuals. For welfare evaluation, this extension requires the choice of a social welfare function, a potentially controversial issue. Sources of controversy include the existence of a social welfare function, ordinal versus cardinal

[2]Details and some caveats are given below.

[3]Two classic examples are Hall (1988) and Campbell and Mankiw (1989) who find the IES to be close to zero.

ordering, and aggregation (see Sen 1970 for a discussion). Fortunately, Lucas' assertion suggests a natural choice for the welfare function. First, Lucas' quantitative claim presumes the existence of a cardinal ordering of preferences that makes interpersonal comparison of utilities possible. Second, he proposes that everyone should be 'equally valued' in the social welfare function. These two observations suggest the use of a classical utilitarian or Benthamite social welfare function, a function that weighs everyone equally.

We further assume, following a standard practice in macroeconomics, that all individuals have identical preferences. This assumption is suggested by Lucas' idea of treating the American and Indian children as 'brothers', and gives further support to the use of an utilitarian welfare function.[4] Our model thus assumes that inequality arises from unequal opportunities rather than differences in tastes. Due to the concavity of the individual's utility function, our social welfare function implies that any inequality is socially costly. By construction, the social welfare function penalizes consumption differences across individuals in the same way that individuals penalize consumption differences across states and time.

We use the extended model to compute welfare costs of within-country and cross-country inequality. Measures of within-country consumption inequality are based on US estimates from Krueger and Perri (2002). Measures of cross-country consumption inequality are based on the Penn World Tables. Inequality in our model is fully characterized by the standard deviation of cross-sectional consumption. The cost of inequality is defined as a permanent compensation rate on consumption required to leave the social planner indifferent between the observed situation and an ideal situation with no cross-sectional dispersion.

We find that the welfare costs of inequality are large and, perhaps surprisingly, likely larger than the gains of economic growth. They range from 12% to 91% for within-country inequality, and from 40% to almost 100% for cross-country inequality, for various values of the IES. Why are the potential welfare costs of inequality seemingly larger than the gains of growth? After all, economic growth affects all individuals alike and therefore has a first-order effect on aggregate welfare. Changes in

[4] Harsanyi (1953, 55, 82), Vickey (1960), and Mirless (1982), among others, support the use a utilitarian welfare function on ethical grounds. Pattanaik (1968) criticizes utilitarianism but recognizes that it is the proper formulation if individuals have identical preferences.

inequality, in contrast, have offsetting effects on individuals and therefore should have only second-order effects on social welfare. This conjecture turns out to be incorrect for three reasons. First, most of the gains of economic growth occur in the future while the costs of inequality are borne every period. Second, inequality is large. While the standard deviation of the log of aggregate consumption is around 1 to 2% — which explains the small welfare costs of business cycles uncovered by Lucas — the standard deviation of the cross sectional distribution of log consumption is around 50% within countries and around 100% across countries. These figures imply that consumption dispersion is larger than mean consumption both within countries and across countries. Third, commonly used values of the IES entail substantial social aversion to any source of consumption dispersion, in particular social aversion to inequality.

An alternative way to understand why inequality is so costly is to interpret the welfare function as the lifetime utility of a newborn child, an interpretation that is probably closer to Lucas' assertion.[5] Under this interpretation, the social planner cares only about newborns and values all children equally, the 'American child and his Indian brother.' If given a choice before birth, what level of growth and inequality would such a child choose behind this Rawlsian 'veil of ignorance' knowing that his place of birth is random? Although growth is clearly welfare-enhancing, this hypothetical child faces huge consumption risk, and a non-trivial probability of poverty, which substantially reduces expected utility.

Further insight into the cost of inequality is obtained by performing counterfactual experiments. We ask: (i) how much growth would society be willing to give up for perfect equality? (ii) what level of inequality would compensate for the lack of economic growth? and (iii) what would be the welfare consequences of eliminating all growth and inequality simultaneously?

Consider the answer to these questions for an intermediate value of the IES of 1/2. In response to the first question we find that a planner would give up 1.62 points of economic growth (out of 2.1 points) to eliminate all within-country inequality, and 4.49 points to eliminate all cross-country inequality. The answer to the second question is that a reduction of 34% of cross-country inequality or 136% of within-country inequality would compensate for the lack of growth. For the third question,

[5] Harsanyi (1953, 1955) and Vickrey (1960) provide similar interpretations of the social welfare function.

we find that eliminating all growth and cross-country inequality is welfare improving — a 48% welfare gain — but eliminating growth and within-country inequality reduces welfare — a 9% welfare fall. Overall, these figures suggest that cross-country inequality is the major determinant of worldwide welfare, and that within-country inequality is as important as growth for country welfare.

An alternative way to evaluate Lucas' assertion is to compute the shadow price of inequality in terms of growth. This shadow price is the marginal willingness to substitute inequality for growth $\left(\frac{\partial\mu}{\partial\sigma} = \frac{\partial \text{growth}}{\partial \text{inequality}}\right)$. If Lucas is right, this price must be close to zero. Moreover, if social decisions are nearly optimal then this shadow price provides not only an assessment of the social willingness to trade inequality for growth, but also an assessment of what is technologically feasible. We find that the price of inequality is typically far from zero. For various values of the IES, this price ranges from 0.026 to 6.55 for within-country inequality, and from 0.0521 to 13.11 for cross-country inequality.

These basic welfare measures provide only upper bounds to the actual gains and costs because they ignore the costs of redistributive policies, and in particular, overlook potential tradeoffs between efficiency and inequality. Specifically, the welfare gain of economic growth ignores the possibility that additional growth may foster inequality. Similarly, the welfare cost of inequality ignores the real possibility that eliminating inequality may also adversely affect incentives, reduce economic growth, and the level of consumption.

Our third extension to Lucas (1987) is to introduce technological constraints. For this purpose, we restrict our analysis to within-country inequality. We postulate a simple reduced form technology that seeks to capture the major tradeoffs that a social planner faces at the country level. The technological frontier is defined in the space of inequality, growth, and consumption levels. We provide two different calibrations of this frontier using US and Scandinavian data from Aaberge et al. (2002), and data from West and East Germany from various sources. The model is then used to compute optimal levels of inequality, growth, and initial consumption level.

We find that the current US values of inequality, growth, and consumption levels are close to their optimal values if the IES is around 1/2. However, if this elasticity is smaller, as various studies suggest, then there is excessive inequality in the US.[6] For example, if the IES is 1/5, then the optimal level of inequality is close to the one observed in Scandinavian countries, and the welfare costs of maintaining current suboptimal choices is around 15%.

In conclusion, we show that the gross and net cost of inequality are likely to be large. Instead of Lucas or Ricardo's, our findings supports Okun's view:

> "I am wandering away from my usual concerns briefly to discuss an even more nagging and pervasive tradeoff, that between inequality and efficiency. It is, in my view, our biggest socioeconomic tradeoff, and it plagues us in dozens of dimensions of social policy." (Okun, 1975, page 2)

This article is organized as follows. Section 2 develops a model for welfare analysis based on Lucas (1987), derives welfare measures, calibrates the parameters, and reports results; Section 3 extends the model of Section 2 by introducing technological restrictions, derives optimal choices, calibrates the parameters, and reports results and Section 4 concludes.

2 Welfare Analysis

2.1 The Model

2.1.1 The Distribution of Consumption

Consider a world composed of a large number of countries (a continuum), with each country equally populated by a large number of individuals (also a continuum). The size of the world population is normalized to 1. The time-t consumption of a particular individual in the world is described by the autoregressive process

$$\ln c_t = \rho \ln c_{t-1} + (1-\rho)(a+bt) + \sigma_\eta \eta_t + \sigma_\epsilon \epsilon_t, \tag{1}$$

[6] A major reason to believe that this elasticity is small is the equity-premium puzzle. See Kocherlakota (1996, page pg 52).

where $\rho \in [0,1)$ is the persistence of consumption, a and b are constants, η_t is a country-specific shock, and ϵ_t is an individual-specific shock, both assumed to be independently drawn from a standard normal distribution. Individuals in the same country share the same draw of η_t. Under these assumptions, the unconditional distribution of consumption satisfies

$$\ln c_t \sim N\left(a + bt, \sigma_x^2 + \sigma_y^2\right), \tag{2}$$

where

$$\sigma_x^2 \equiv \frac{\sigma_\eta^2}{1-\rho^2};\ \sigma_y^2 \equiv \frac{\sigma_\epsilon^2}{1-\rho^2}. \tag{3}$$

We assume that the law of large number holds so that there is no aggregate uncertainty at the worldwide level. Furthermore, we assume that the initial distribution of consumption is given by its unconditional distribution evaluated at time $t = 0$:

$$\ln c_0 \sim N\left(a, \sigma_x^2 + \sigma_y^2\right). \tag{4}$$

These assumptions imply that the worldwide distribution of consumption at any point in time is described by (2). According to this cross-sectional interpretation of (2), σ_y^2 measures the degree of within-country inequality — inequality associated with individual factors — and σ_x^2 measures cross-country inequality — inequality associated with country-specific factors. This framework can be used to analyze a single country if $\sigma_x^2 = 0$, or a group of countries populated by identical individuals if $\sigma_y^2 = 0$.

Parameters a and b are chosen so that aggregate worldwide consumption at time t, Ec_t, equals $(1+\lambda)(1+\mu)^t$. This requires to set $a = \ln(1+\lambda) - \frac{1}{2}\left(\sigma_x^2 + \sigma_y^2\right)$ and $b = \ln(1+\mu)$. In this specification, μ is the growth rate of worldwide consumption, and λ is the worldwide level of consumption at time 0, which determines the level of consumption in any subsequent period. λ is used below to measure the welfare gains and costs of different experiments, and it is set to 0 in the baseline case.

2.1.2 Individual and Social Welfare

The welfare of an individual with initial consumption c_0 is described by the expected utility function

$$U(c_0) = E_0 \left[\sum_{t=0}^{\infty} \beta^t u(c_t) \right], \tag{5}$$

where E_0 is the mathematical expectation conditional on time-0 information, and u is a momentary utility function, identical for all individuals. We define social welfare as the average welfare of the society,

$$W \equiv \int U(c) F(c) dc, \tag{6}$$

where $F(c)$ is the fraction of the population with consumption below or equal to c at time 0. F is implicitly defined by (4). There are least two alternative interpretations of W. First, W is a standard utilitarian social welfare function that weighs everyone's welfare equally. Second, W is the expected welfare of a newborn child, EU. Furthermore, depending on the values of σ_x and σ_y, a 'society' could designate the world — if $\sigma_x > 0$ and $\sigma_y > 0-$ or a particular country — if $\sigma_x = 0$. Our choice of the social welfare function is inspired by Lucas' (2004) suggestion that a social planner ('God') places equal weight ('equally valued') on individuals similar at birth ('American and Indian brothers'). The difference between individuals is not to be found in different tastes, but different endowment of resources ('The resources of all kinds that will be at the disposal of this new American will be on the order of 15 times the resources available to his Indian brother.').

We further simplify the problem by assuming $u(c) = \frac{c^{1-\gamma}-1}{1-\gamma}$, where $1/\gamma > 0$ is the intertemporal elasticity of substitution (IES), and γ is the coefficient of relative risk aversion. Given that F is given by (4), W satisfies:

$$\begin{aligned} W &= EU = EE_0 \left[\sum_{t=0}^{\infty} \beta^t u(c_t) \right] = \sum_{t=0}^{\infty} \beta^t \frac{Ec_t^{1-\gamma} - 1}{1-\gamma} \\ &= \frac{1}{1-\gamma} \sum_{t=0}^{\infty} \beta^t \left(e^{(1-\gamma)(a+bt)+\frac{1}{2}(1-\gamma)^2\left(\sigma_x^2+\sigma_y^2\right)} - 1 \right). \end{aligned}$$

Substituting for the definitions of a and b, and simplifying produces

$$W = W(\lambda, \mu, \sigma_x^2 + \sigma_y^2; \gamma, \beta) \equiv \frac{(1+\lambda)^{1-\gamma} e^{-\gamma(1-\gamma)\left(\sigma_x^2+\sigma_y^2\right)/2}}{(1-\gamma)\left(1-\beta(1+\mu)^{1-\gamma}\right)} - \frac{1}{1-\gamma}\frac{1}{1-\beta} \tag{7}$$

Accordingly, social welfare depends on the initial level of consumption, λ, the growth of consumption, μ, the total dispersion of consumption, measured by $\sigma_x^2 + \sigma_y^2$, and preference parameters γ and β. It easy to check that inequality reduces social welfare ($\frac{\partial W}{\partial \sigma^2} < 0$) and growth increases social welfare ($\frac{\partial W}{\partial \mu} > 0$). Notice that the degree of persistence, ρ, only affects social welfare through its effect on consumption dispersion, $\sigma_x^2 + \sigma_y^2$. Thus, social mobility only matters for social welfare to the extent that it affects the level of inequality. Furthermore, the social welfare function implicitly assumes that the coefficient of inequality aversion equals the coefficient of risk aversion, and the inverse of the intertemporal elasticity of substitution (Atkinson 1970).[7] In other words, the planner penalizes consumption dispersion across individuals in the same way that individuals penalize consumption dispersion across states and time.

In order to perform welfare comparisons, we need to compute a baseline welfare level, W_0, determined by a baseline set of parameters $[\mu_0, \sigma_{x0}^2, \sigma_{y0}^2]$. Let $\sigma_0^2 \equiv \sigma_{x0}^2 + \sigma_{y0}^2$. The baseline welfare level satisfies

$$W_0 = W(0, \mu_0, \sigma_0^2) \tag{8}$$

The following subsections use the above framework to define a number of welfare measures. We first define welfare costs and gains as the proportional increase in consumption necessary to leave the social planner indifferent between a baseline and an alternative situation. We then present the welfare costs of inequality as the amount of growth a planner would give up to eliminate inequality. Finally, we present some preliminary measures of the tradeoff between inequality and growth — first, by quantifying the welfare costs of eliminating all growth and all inequality simultaneously; second in the form of a marginal rate of substitution between inequality and growth.

[7] See also Auerbach and Hassett (1999), and Kaplow (2003) for a discussion and generalizations, and Harsanyi (1975) for a critique of the generalizations.

2.1.3 Standard Welfare Measures

Following Lucas (1987), we define four measures of social gains (or costs, if negative), λ_μ, λ_x, λ_y and λ_0, as solutions to the following equations:

$$\begin{aligned} W_0 &= W(\lambda_\mu, 0, \sigma_0^2), \\ W_0 &= W(\lambda_x, \mu_0, \sigma_0^2 - \sigma_{x0}^2), \\ W_0 &= W(\lambda_y, \mu_0, \sigma_0^2 - \sigma_{y0}^2). \\ W_0 &= W(\lambda_0, \mu_0, 0). \end{aligned}$$

These values of $\lambda's$ are the proportional changes of consumption required to leave the world planner indifferent between the baseline consumption path with welfare W_0, and an alternative consumption path with welfare W. Notice that the $\lambda's$ affect the consumption path of the alternative situation, not the baseline situation. Thus, for example, λ_μ is the proportional change in consumption, uniform across all periods, countries, and individuals, required to leave the world planner indifferent between the baseline consumption path, and a path with no growth (but a higher initial consumption level measured by λ_μ). Notice also that $\lambda_\mu > 0$ and $(\lambda_x, \lambda_y, \lambda_0) < 0$. These measures can be interpreted in the following way: λ_μ is the welfare *gain* of economic growth, λ_x is the welfare *cost* of cross-country inequality, λ_y is the welfare *cost* of within-country inequality, and λ_0 is welfare cost of total inequality. The equations above have the following simple solutions:

(9) $$\lambda_\mu = \left[\frac{1-\beta}{1-\beta\left(1+\mu_0\right)^{(1-\gamma)}}\right]^{\frac{1}{1-\gamma}} - 1,$$

and

(10) $$\lambda_i = e^{-\gamma\sigma_i^2/2} - 1 \text{ for } i = \{x, y, 0\}.$$

These formulas reveal three important properties of the welfare measures. First, welfare measures depend only on the single relevant parameter. For example, λ_y depends only on σ_y^2 but not on μ or σ_x^2.

This is mainly a consequence of assuming an isoelastic utility function and a log-normal distribution of consumption. An important implication of this result is that the welfare costs of within-country inequality and gains of economic growth are the same regardless of whether the society is the world (so that $\sigma_x > 0$) or a country (so that $\sigma_x = 0$).

Second, λ_μ is strictly decreasing in γ : growth is less attractive if consumers are less willing to substitute consumption intertemporally, a result that is well-known in the risk-free rate puzzle literature (see Kocherlakota 1996). Third, λ_i increases exponentially with γ. This result follows naturally from the fact that a more concave momentary utility function makes any dispersion of consumption more costly. These last two results imply that the welfare gains of economic growth can be made arbitrarily small and the welfare costs of inequality arbitrarily large by increasing γ.

Another interesting welfare measure is the gain associated with one additional percentage point of economic growth, $\lambda_{1\%}$. It is defined as

$$W_0 = W(\lambda_{1\%}, \mu_0 - 0.01, \sigma_0^2).$$

Using (7), $\lambda_{1\%}$ is given by

$$\lambda_{1\%} = \left[\frac{1 - \beta \left(1 + \mu_0 - 0.01\right)^{(1-\gamma)}}{1 - \beta \left(1 + \mu_0\right)^{(1-\gamma)}} \right]^{\frac{1}{1-\gamma}} - 1. \tag{11}$$

2.1.4 Alternative Welfare Measures

The welfare measures just defined cannot be easily compared to each other. For example, λ_μ is a compensation rate on a flat consumption path while λ_y is a compensation rate on a growing consumption path. Thus, for example, a case in which $\lambda_\mu = -\lambda_y$ would actually imply that the welfare gains of economic growth are lower than the welfare costs of within-country inequality. The following welfare measures have the advantage of being defined as compensation rates on the same consumption path, the baseline path. Define $\widehat{\lambda}_\mu$, $\widehat{\lambda}_x$, $\widehat{\lambda}_y$ and $\widehat{\lambda}_0$ as the solutions to the following equations:

$$
\begin{aligned}
W(\widehat{\lambda}_\mu, \mu_0, \sigma_0^2) &= W(0, 0, \sigma_0^2), \\
W(\widehat{\lambda}_x, \mu_0, \sigma_0^2) &= W(0, \mu_0, \sigma_0^2 - \sigma_{x0}^2), \\
W(\widehat{\lambda}_y, \mu_0, \sigma_0^2) &= W(0, \mu_0, \sigma_0^2 - \sigma_{y0}^2), \\
W(\widehat{\lambda}_0, \mu_0, \sigma_0^2) &= W(0, \mu_0, 0).
\end{aligned}
$$

The $\widehat{\lambda}'s$ are thus defined as percentage changes relative to the baseline consumption path rather than the alternative consumption path. These definitions imply that $\widehat{\lambda}_\mu < 0$ and $(\widehat{\lambda}_x, \widehat{\lambda}_y, \widehat{\lambda}_0) > 0$. These measures can be interpreted in the following way: $\widehat{\lambda}_\mu$ is the welfare *cost* of no growth, $\widehat{\lambda}_y$ is the welfare *gain* of eliminating within-country inequality, $\widehat{\lambda}_x$ is the welfare *gain* of eliminating cross-country inequality, and $\widehat{\lambda}_0$ is the welfare gain of eliminating all inequality. The equations above have the following simple solutions:

(12)
$$\widehat{\lambda}_\mu = \frac{1}{1+\lambda_\mu} - 1,$$

and

(13)
$$\widehat{\lambda}_i = \frac{1}{1+\lambda_i} - 1 \text{ for } i = \{x, y, 0\}.$$

and similarly for $\widehat{\lambda}_{1\%}$.

2.1.5 Welfare Measures in Terms of Growth Rates

Another way to compensate the planner for alternative consumption paths is through changes in the growth rate of consumption, μ, rather than through changes in the level of consumption, λ. One can define welfare cost measures in terms of growth rates, μ_x, μ_y, and μ_0 as follows:

$$\begin{aligned} W_0 &= W(0, \mu_0 + \mu_x, \sigma_0^2 - \sigma_{x0}^2), \\ W_0 &= W(0, \mu_0 + \mu_y, \sigma_0^2 - \sigma_{y0}^2), \\ W_0 &= W(0, \mu_0 + \mu_0, 0). \end{aligned}$$

According to these definitions, μ_y represents the additional points of economic growth that the planner would be willing to accept in exchange for eliminating within-country inequality. Naturally, $\mu_y < 0$. Similarly, μ_x and μ_0 are the points of economic growth that the planner would be willing to accept in exchange for eliminating cross-country and total consumption inequality respectively. Using (7) and the definitions above, these values of μ satisfy:

$$\mu_i = \left\{ \frac{1}{\beta} \left[1 - e^{\gamma(1-\gamma)\sigma_i^2/2} \left(1 - \beta \left(1 + \mu_0\right)^{(1-\gamma)} \right) \right] \right\}^{\frac{1}{1-\gamma}} - (1 + \mu_0) \text{ for } i = \{x, y, 0\}. \tag{14}$$

2.1.6 Welfare Gains of Growth in Terms of Inequality

An alternative way to compensate the planner for the absence of economic growth is through changes in the levels of inequality rather than in the levels of consumption. For this purpose, define the rates θ_x, θ_y, and θ_0 as follow:

$$\begin{aligned} W_0 &= W(0, 0, \sigma_0^2 - \theta_x \sigma_{x0}^2), \\ W_0 &= W(0, 0, \sigma_0^2 - \theta_y \sigma_{y0}^2), \\ W_0 &= W(0, 0, \sigma_0^2 - \theta_0 \sigma_0^2). \end{aligned}$$

According to these definitions, $\theta_i \times 100$ is the percentage reduction in inequality (either cross-country, within-country, or total inequality) that a planner would be willing to accept in exchange for zero economic growth. Using (7), these measures can be computed as:

$$\theta_i = \frac{2 \left[\ln (1 - \beta) - \ln \left(1 - \beta \left(1 + \mu_0\right)^{(1-\gamma)} \right) \right]}{\gamma (1 - \gamma) \sigma_i^2} \text{ for } i = \{x, y, 0\}. \tag{15}$$

2.1.7 Net Welfare Consequences of Growth and Inequality

It is often argued in the literature that inequality is partly the result of providing individuals with incentives conducive to economic growth. It is therefore natural to wonder what the consequences of eliminating inequality and economic growth at the same time would be. The welfare consequences of these experiments are given by $\overline{\lambda}_x$, $\overline{\lambda}_y$, and $\overline{\lambda}_0$ defined as:

$$\begin{aligned} W_0 &= W(\overline{\lambda}_x, 0, \sigma_0^2 - \sigma_{x0}^2), \\ W_0 &= W(\overline{\lambda}_y, 0, \sigma_0^2 - \sigma_{y0}^2), \\ W_0 &= W(\overline{\lambda}_0, 0, 0), \end{aligned}$$

Thus, $\overline{\lambda}$ is a *net* welfare gain (cost if negative) of the current inequality-growth combination. Using the definitions above and (7), it follows that:

$$\overline{\lambda}_i = e^{-\gamma\sigma_i^2/2}\left[\frac{1-\beta}{1-\beta\left(1+\mu_0\right)^{(1-\gamma)}}\right]^{\frac{1}{1-\gamma}} - 1 \text{ for } i = \{x, y, 0\} \tag{16}$$

2.1.8 The Shadow Price of Inequality

One way to interpret Lucas' assertion is that a society must be willing to trade only a small reduction in growth for a one point reduction in inequality. If so, the social marginal rate of substitution between inequality and growth must be close to zero. We can compute this marginal rate around (σ_{i0}, μ_0) using (7), holding $W = W_0$ and $\lambda = 0$.

We find that

$$MRS_1^i = \frac{\partial \mu}{\partial \sigma_i} = \sigma_{i0}\gamma\frac{\left(1+\mu_0\right)^{\gamma} - \beta\left(1+\mu_0\right)}{\beta} \text{ for } i = \{x, y, 0\}$$

Lucas' assertion suggests $MRS_1^i = 0$. Notice that MRS_1^i increases linearly with the degree of inequality and exponentially with γ.

One can alternatively compute the marginal rate of substitution between inequality and consumption level, rather than growth. This shadow price is given by:

$$MRS_2^i = \frac{\partial \lambda}{\partial \sigma_i} = \sigma_{i0}\gamma \text{ for } i = \{x, y, 0\}\,.$$

2.2 Calibration

In order to compute welfare measures we need to calibrate the parameters μ_0, σ_{x0}^2, σ_{y0}^2, γ and β. Figure 1 shows the unweighted and weighted averages of the log of consumption per capita in 108 economies between 1960 and 2000. The data is from the Penn World Tables 6.1. Population size from each country is used to compute weighted averages. The average growth rate of yearly per-capita consumption is 2.3% and its standard deviation is 1.17% for the unweighted average, and 2.1% and 0.99% respectively for the weighted average. This evidence suggests that $\mu_0 = 2.1\%$.

Figure 2 shows the unweighted and weighted standard deviations of the log of per-capita consumption across countries from 1960 to 2000. The unweighted and weighted averages provide different characterizations of the evolution of cross-country dispersion. The unweighted average suggests a significant increase of consumption dispersion during this 40-year period. However, the weighted average suggests that the dispersion has remained roughly constant over the long term, and in fact, has actually decreased during the past 25 years, after increasing significantly during the 1960-1975 period. Figure 2 thus suggests that $\sigma_{x0} = 1$.

International evidence about the dispersion of consumption within countries is scarce. Some authors report measures of income dispersion within countries (e.g. Bourguignon and Morrison (2002), Sala-i-Martin (2002)), but not of consumption dispersion. Differences in dispersion between consumption and income may be significant. Krueger and Perri (2002, Figure 1) provide some estimates about the dispersion of individual consumption and income for the United States. They find that the standard deviation of the log of individual consumption, controlling for age and race, has been roughly constant in a 25-year period at around 0.48. The dispersion of per-capita income is about twice as large. Based on this evidence, we choose $\sigma_{y0} = 0.5$.

Numerous studies present an estimate of γ but there is little consensus about its proper value. At one extreme, Hall (1988) and Campbell and Mankiw (1989) find the IES to be close to zero (or

$\gamma \simeq \infty$).[8] At the other extreme, Beaudry and van Wincoop (1996) suggest that the IES is close to 1 ($\gamma \simeq 1$). We follow Lucas (1987) and compute welfare measures for different values of the IES $\frac{1}{\gamma}$ or for $\gamma \in [1, 2, 5, 10, 20]$.[9]

Finally, we calibrate β using the standard procedure of targeting a value for the risk free interest rate. This implies that β depends on γ. Since there is country risk in our framework, only bonds traded internationally can be completely risk free. If r is the interest rate paid on a risk free bond, individual optimization implies the following Euler equation:

$$c_t^{-\gamma} = \beta (1+r) E_t c_{t+1}^{-\gamma}.$$

According to equation (1), $c_{t+1} = c_t^{\rho} e^{(a+b(t+1))(1-\rho)} z_{t+1}$, where $\ln z_{t+1} = \sigma_\eta \eta_t + \sigma_\epsilon \epsilon_t$. Thus,

$$c_t^{-\gamma} = \beta (1+r) c_t^{-\gamma\rho} e^{-\gamma(a+b(t+1))(1-\rho)} e^{\frac{\gamma^2}{2}\left(\sigma_\eta^2+\sigma_\epsilon^2\right)},$$

and solving for c_t produces

$$c_t = e^{a+b(t+1)} e^{-\frac{\gamma}{2}\frac{\sigma_\eta^2+\sigma_\epsilon^2}{1-\rho}} \left[\beta (1+r)\right]^{-\frac{1}{\gamma(1-\rho)}}.$$

Aggregating across agents, using (2), implies

$$Ec_t = e^{a+bt+\frac{1}{2}\frac{\sigma_\eta^2+\sigma_\epsilon^2}{1-\rho}} = e^{a+b(t+1)} e^{-\frac{\gamma}{2}\frac{\sigma_\eta^2+\sigma_\epsilon^2}{1-\rho}} \left[\beta (1+r)\right]^{-\frac{1}{\gamma(1-\rho)}}.$$

Finally, simplifying and solving for $1+r$ produces:

$$1+r = \underbrace{\frac{(1+\mu)^\gamma}{\beta}}_{\text{Standard Effect}} \underbrace{(1+\mu)^{-\rho\gamma}}_{\text{persistence}} \underbrace{e^{-\frac{\gamma(1+\gamma)}{2}\left(\sigma_\eta^2+\sigma_\epsilon^2\right)}}_{\text{risk}}.$$

Note the different components of this expression. The first is a standard purely deterministic com-

[8] Many researchers regard the equity premium puzzle as evidence that γ is large. See Kocherlakota (1996, pg. 52) for references.

[9] We approximate the log case ($\gamma = 1$) by using $\gamma = 1.01$ in the formulas above instead of providing additional derivations.

ponent, inducing a positive relationship between γ and r. This positive relationship is the source of the risk-free rate puzzle (see Kocherlakota 1996). For a given level of r, this first component would imply a positive relationship between γ and β. The second component captures the persistence of consumption overtime, an effect that drives the risk-free rate downward as γ increases. The final component captures precautionary motives. The presence of individual and country risks induce precautionary savings that drive the risk-free rate downward as γ increases. For a given target level of r, the two last components induce a negative relationship between β and γ.

The first effect dominates only if consumption is not very persistent and consumption risk is small. However, according to equation (3) the observed large degree of within-country and cross-country inequality must come either from significant consumption persistence or from large consumption risk, or both. The last two effects turn out to dominate producing a negative relationship between β and γ. Using (3) into the last equation and solving for β gives:

$$\beta = \frac{(1+\mu)^{\gamma(1-\rho)}}{(1+r)} e^{-\frac{\gamma(1+\gamma)\left(1-\rho^2\right)}{2}\left(\sigma_x^2+\sigma_y^2\right)}.$$

We assume that the model's period is one year and pick $r = 5\%$. This choice allows us to replicate $\beta = 0.95$ for $\gamma = 1$, the value used by Lucas (1987).[10] We find that β decreases dramatically with γ unless ρ is very close to 1. To prevent a dramatic fall of β for large γ we set $\rho = 0.9999$.[11] This choice implies relative small consumption risk, $\sigma_\eta = 1.41\%$ and $\sigma_\epsilon = 0.71\%$, and a β equal to 0.90 for $\gamma = 20$. Lower values of ρ will increase annual risk and strengthen our results.[12] Recall that $\sigma_\eta^2 = \left(1-\rho^2\right)\sigma_{x0}^2$ and $\sigma_\epsilon^2 = \left(1-\rho^2\right)\sigma_{y0}^2$. Given the calibrated values of σ_{x0} and σ_{y0}, the choice of ρ only affects the values of σ_η and σ_ϵ.

All parameter values are summarized in Table 1.

[10] A value for r around 5% is typically used in the literature. For example, King and Rebelo (2000) use $r = 6.5\%$ which is the postwar U.S. average of the annual return on capital. A larger value of r will strengthen our results.

[11] A large ρ is consistent with Constantiniades and Duffie (1995), Gourinchas and Parker (2002), Storeletten et al. (2004), Castañeda et al. (2003), but inconsistent with Heaton and Lucas (1996), and Aiyagari (1994).

[12] We could alternatively calibrate the model for a version without aggregate risk at the country level. In that case $\beta = \frac{(1+\mu)^{\gamma(1-\rho)}}{(1+r)} e^{-\frac{\gamma(1+\gamma)\left(1-\rho^2\right)}{2}\sigma_y^2}$. The large value of σ_y^2 still implies a negative relationship between β and γ. Moreover, we would still need a large value of ρ (=0.9995) to obtain $\beta = 0.9$ for $\gamma = 20$.

2.3 Results

Table 2 reports the welfare measures defined in Section 2.2.1 for the parameters described above and for different values of γ and β. Recall that λ_μ is the welfare gain of economic growth, λ_x is the welfare cost of cross-country inequality, λ_y is the welfare cost of within-country inequality, and λ_0 is the welfare cost of total inequality. The table reproduces Lucas' (1987) finding that the welfare gains from economic growth, as described by λ_μ and $\lambda_{1\%}$, are substantial if γ is close to 1. These gains, however, significantly decrease as γ increases. For example, Lucas reports a 20% gain from 1 additional point of economic growth, but the table shows that if γ takes an intermediate value of 5, the gains are 8.5% instead. Total gains of economic growth as measured by μ are between 8% and 51%.

The novel and surprising result reported in Table 2 is the large welfare costs associated with consumption inequality, both within and across countries. Costs of inequality range from around 12% to 92% for within-country inequality, and from around 40% to almost 100% for cross-country inequality. In fact, the gains of economic growth are smaller than the cost of total inequality if $\gamma \geq 1.11$, smaller than the cost of cross-country inequality if $\gamma \geq 1.28$, and smaller than the cost of within-country inequality if $\gamma \geq 3.2$.

Overall, the results reported in Table 2 contradicts Lucas' (2004) assertion that the benefits of growth dwarf the costs of inequality. These figures are staggering and reveal how important inequality is for aggregate welfare. In recent years, many authors have argued that Lucas' original calculation understates the true cost of business cycles because of his restrictive assumptions on preferences, homogeneity of consumers, or his choice of matching the risk implied in aggregate rather than individual data (see Lucas 2003 or Barlevy 2005 for a survey of this literature). When some of these assumptions are relaxed, the cost of business cycle can reach 3-4% for households with no wealth (e.g. Beaudry and Pages 2001, or Krusell and Smith 1999) or be as large as 12% with less restrictive preferences (e.g. Tallarini 2000). But even these estimates of the costs of business cycles seem small relative to the costs of inequality.

Table 3 reports the alternative welfare measures described in Section 2.2.2. Recall that $\widehat{\lambda}_\mu$ is the welfare cost of no growth, and $\widehat{\lambda}_x$ is the welfare gain of eliminating cross-country inequality, $\widehat{\lambda}_y$ is the welfare gain of eliminating within-country inequality. The message is similar to the one in Table 2, but the magnitudes, particularly regarding the welfare consequences of inequality, are even more staggering. These welfare measures can be interpreted as proportional taxes (if negative) or proportional subsidies (if positive) on the baseline consumption path that leaves the planner indifferent between the after-tax (subsidy) baseline consumption path and the alternative path. Consider $\gamma = 2$ for example. Eliminating economic growth is equivalent to introducing a permanent 28% tax on consumption, eliminating within-country inequality is equivalent to introducing a permanent 28% subsidy on consumption, and eliminating cross-country inequality is equivalent to introducing a permanent 249% subsidy on consumption (!).

The major drawback of these welfare measures is that they ignore potential trade-offs between efficiency and inequality. Higher economic growth may cause more inequality. On the other hand, inequality may foster growth. For example, the prospect of appropriating higher returns without the fear of redistributive taxation may spur investment, effort, and innovation. The next section addresses this concern by introducing a technological tradeoff between inequality and consumption growth and levels. In this section, we provide an indirect assessment by studying the planner's willingness to substitute equality for growth.

A first measure of this tradeoff is given by the reduction in growth that would perfectly offset the gains of eliminating inequality. These welfare measures were described in Section 2.2.3 and are reported in Table 4 under the labels μ_i, $i = \{x, y, 0\}$. For example, $\mu_y = -3.1\%$ for $\gamma = 5$ means that the planner would accept a reduction of 3.1 points in the growth rate of consumption in exchange for eliminating all within-country inequality. Since the baseline growth rate is 2.1%, the new growth rate would be -1%. All the $\mu-$welfare measures are negative, in most cases implying negative net growth rates, which again reveals the staggering social costs of inequality.

A second measure of this tradeoff is given by the reduction in inequality that would compensate the planner for the elimination of growth. These welfare measures were described in Section 2.2.4 and are reported in Table 4 under the labels θ_i, $i = \{x, y, 0\}$. For example, $\theta_y = 36\%$ for $\gamma = 5$ means that the planner only requires a 36% reduction of within-country inequality to be compensated for the lack of growth. The fact that most $\theta' s$ are below 36% suggests that relative small reductions in inequality would compensate for the absence of growth.

A third measure of this tradeoff are the welfare consequences of eliminating growth and inequality simultaneously. These welfare measures were described in Section 2.2.5 and are reported in Table 4 under the labels $\overline{\lambda}_i$, $i = \{x, y, 0\}$. For example, $\overline{\lambda}_y = -33\%$ for $\gamma = 5$ means a large welfare gain of eliminating inequality and growth. The fact that most $\overline{\lambda_i}'$ are negative suggests that there is too much inequality relative to growth.[13]

A final measure of this tradeoff is given by the social marginal rate of substitution between inequality and growth, defined in Section 2.2.6. and reported in Table 4 under the label MRS_1^i for $i = \{x, y, 0\}$. Given our framework, we can construct social indifference curves between inequality and growth defined by $W(0, \mu, \sigma^2) = W_0$. [14] This indifference curve traces all combinations of growth rates and consumption dispersion of alternative consumption paths that produce the same welfare as the baseline situation. The slope of this map determines the willingness of the planner to trade inequality for growth. The top panel of Figure 3 illustrates these indifference maps for various values of γ. According to Lucas (2004), the shadow price of inequality must be close to zero, or alternatively the social indifference curve must be flat. The table shows, however, that this is not generally the case. For example, for $\gamma = 5$, $MRS_1^y = \frac{\partial \mu}{\partial \sigma_y} = 0.371$ means that the shadow price of 1 point of inequality, measured by σ, is around $\frac{1}{3}$ of a point of economic growth. Even for $\gamma = 1$ the shadow price of inequality is significantly different from zero.

Finally, Table 4 also reports the marginal rate of substitution between inequality and the consumption level (MRS_2^i) defined in Section 2.2.6. This is a more relevant tradeoff, as we discuss in the next Section, because growth rates tend to be similar across countries in spite of different degrees of inequality, but consumption levels vary widely. This shadow price of inequality is also different from zero, and in most cases significantly larger than one. For example, $MRS_2^y = \frac{\partial \lambda}{\partial \sigma_y} = 2.5$ for $\gamma = 5$ means that the planner would be willing to permanently give up 2.5% of consumption for a

[13] For $\gamma \leq 2.5$, eliminating only within-country inequality and growth is detrimental to welfare. However, even in that case a reduction of inequality has a significant first-order effect on welfare. For example, for $\gamma = 1$ eliminating growth would imply a loss of 51%, but also eliminating within-country inequality reduces this loss to 33%.

[14] Solving for this indifference map produces $\sigma^2 = \frac{2}{\gamma(1-\gamma)} \ln \left[\frac{1-\beta(1+\mu_0)^{(1-\gamma)}}{1-\beta(1+\mu)^{(1-\gamma)}} \right] + \sigma_0^2$.

permanent reduction of 1 point of inequality.

Why is inequality so costly? One way to explain the magnitude of the cost is to interpret the welfare function $W(\lambda, \mu, \sigma^2)$ as the lifetime utility of a newborn child. If given a choice before birth, what level of growth and cross-country inequality would such a child choose behind this Rawlsian 'veil of ignorance' knowing that his country of birth is random? Although growth is clearly welfare-enhancing, this hypothetical child faces a non-trivial probability of poverty, which reduces expected utility. As risk aversion rises, increases in inequality must be compensated with much higher growth.

The last two welfare measures also reflect the technological tradeoff if the observed social choices are optimal. In that case, observed choices lie both on the social indifference curve and on the production possibility frontier. This is illustrated in the bottom panel of figure 3. If social choices are suboptimal however, we need to specify technological constraints — i.e. we need to specify what levels of growth, inequality and consumption are actually feasible — in order to determine *actual* welfare costs and gains rather than just *potential* costs and gains. The remainder of the paper addresses this issue.

3 Optimal Inequality

3.1 The Planner's Problem

In this Section we seek to provide some assessment about the optimal level of inequality. For this purpose, we introduce technological restrictions into the welfare framework of the previous section. We assume that social preferences over the consumption level, $1 + \lambda$, the consumption growth rate μ, and the inequality level, σ^2, are described by the social welfare function (7). For convenience, we rewrite this function dropping the constant,

$$W \equiv \frac{(1+\lambda)^{1-\gamma}\, e^{-\gamma(1-\gamma)\sigma^2/2}}{(1-\gamma)\left(1-\beta\,(1+\mu)^{1-\gamma}\right)}. \tag{17}$$

Furthermore, we postulate that the technological frontier is described by the function

$$\sigma = A\left(1+\mu\right)^{\varepsilon_1}\left(1+\lambda\right)^{\varepsilon_2},\ A>0, \varepsilon_1 \geq 0, \varepsilon_2 \geq 0. \tag{18}$$

This restriction describes the technological tradeoffs between inequality, growth, and consumption levels. This formulation is a reduced form approximation of what is likely to be a very complex determination of consumption levels, growth, and inequality. One can think, for example, that within-country inequality is the result of providing proper incentives to individuals who possess private information about their own abilities, as in Atkenson and Lucas (1992). The social optimal policy may induce significant inequality in order to elicit private information. A full solution of such a model is beyond the scope of this paper.[15] We only provide some suggestive results.

For cross-country inequality, the reason for a technological trade-off between inequality, growth and levels is less clear. Private information issues that could explain the need for within-country inequality do not apply at the country level. When information is public, a worldwide planner could in principle design a non-distortionary tax scheme. The major technological restriction at the worldwide level seems to be an enforcement problem. A worldwide social planner would be unable to enforce any progressive cross-country tax schemes.[16] If no worldwide redistribution is feasible, then current cross-country inequality is the constrained efficient and not much more could be said. None of the large potential welfare gains documented in the previous section can be realized. Since we lack any knowledge about the shape of (18) for cross-country inequality, we focus the remaining analysis on within-country inequality.

[15] Solving a version of such a model would require us to address major theoretical and computational issues. One major problem is getting around the immizeration result to obtain a stationary distribution (see Sleet and Yeltekin 2005, Phelan 2003, and Farhi and Wening 2005 for recent contributions).

[16] A Pareto improvement could be obtained by allowing countries to share their aggregate risk so that consumption at the country level will grow deterministically. Since consumption is highly persistent, the resulting cross-country dispersion of consumption with perfect risk-sharing would be similar to the current distribution. Thus, the major gains from a worldwide planner's perspective would not come from sharing country risk but from redistributing initial resources.

3.2 Optimal Choices

The planner's problem is to choose (σ, μ, λ) to maximize W subject to (18). Substituting (18) into (17) produces

$$V(\lambda, \mu) = \frac{(1+\lambda)^{1-\gamma} e^{-\gamma(1-\gamma)A^2(1+\mu)^{2\varepsilon_1}(1+\lambda)^{2\varepsilon_2}/2}}{(1-\gamma)\left(1-\beta(1+\mu)^{1-\gamma}\right)}.$$

Maximizing this function over λ and μ respectively provides the following optimality conditions:

(19) $$1 = \varepsilon_2 \gamma A^2 (1+\mu^*)^{2\varepsilon_1} (1+\lambda^*)^{2\varepsilon_2} = \varepsilon_2 \gamma \sigma^{*2}$$

(20) $$\frac{\beta(1+\mu^*)^{1-\gamma}}{1-\beta(1+\mu^*)^{1-\gamma}} = \varepsilon_1 \gamma A^2 (1+\mu^*)^{2\varepsilon_1} (1+\lambda^*)^{2\varepsilon_2} = \varepsilon_1 \gamma \sigma^{*2}$$

The first condition determines the optimal degree of inequality, given by

(21) $$\sigma^{*2} = \frac{1}{\varepsilon_2 \gamma}.$$

This expression states that optimal inequality depends positively on the intertemporal elasticity of substitution, $\frac{1}{\gamma}$, and negatively on the elasticity of inequality to the level of consumption. The discount factor and, in particular, ε_1 play no role in determining optimal inequality. Substituting (21) into (20) and solving for $1+\mu^*$ produces:

(22) $$1+\mu^* = \left[\beta\left(1+\frac{\varepsilon_2}{\varepsilon_1}\right)\right]^{\frac{1}{\gamma-1}}$$

This expression states that if $\gamma > 1$ — the empirically relevant case — then the optimal growth rate depends positively on β and ε_2 and negatively on ε_1. Moreover, $\beta\left(1+\frac{\varepsilon_2}{\varepsilon_1}\right) > 1$ is required for positive growth. If this condition is satisfied, then the optimal growth rate depends negatively on γ. Thus, the growth rate is larger if a society is more patient and less risk averse.

Moreover, the optimal consumption level λ^* is found by substituting (21) and (22) into (18) and solving for λ^*. The solution is:

$$(23) \qquad 1+\lambda^* = \left[\frac{\sigma^*}{(1+\mu^*)^{\varepsilon_1}\, A}\right]^{\frac{1}{\varepsilon_2}}.$$

Finally, optimal social welfare is given by $W^* = W(\lambda^*, \mu^*, \sigma^*)$. We can compute the welfare cost associated with the baseline choices, λ^c, as the solution to $W(0, \mu_0, \sigma^2) = (1+\lambda^c)^{1-\gamma}\, W^*$. Thus,

$$(24) \qquad 1+\lambda^c = \left[\frac{W(0, \mu_0, \sigma^2)}{W^*}\right]^{\frac{1}{1-\gamma}}$$

3.3 Calibration

In order to compute optimal inequality we only need an estimate of ε_2. It follows from (18) that $\varepsilon_2 = \frac{\partial \ln \sigma}{\partial \ln(1+\lambda)}$. This elasticity could be computed using the percentage differences in inequality relative to percentage difference in consumption per-capita for two similar countries with similar growth rates. A case that has been well documented is that of the U.S. versus Scandinavian countries (Denmark, Norway, and Sweden). According to the Penn World Tables 6.1, average per-capita income in the U.S. was around 24% higher than the average of Scandinavian countries in 1990, and around 25% higher in 2000. The relative stability of this difference means that the U.S. and Scandinavian countries share similar growth rates. These figures suggest to pick $\lambda = 25\%$ or $\triangle \ln(1+\lambda) = \ln 1.25$.

Regarding differences in inequality, Aaberge et al. (2000, Table 2) report Gini coefficients of disposable income distribution for the US and Scandinavian countries for different years up to 1990. They find a Gini coefficient of 0.346 for the U.S. and an average Gini of 0.2173 for Scandinavian in 1990. Assuming that disposable income is lognormal distributed, Gini coefficients can be transformed into standard deviations using the formula $Gini = 2\Phi\left(\frac{\sigma}{\sqrt{2}}\right) - 1$where Φ is the standard normal distribution. Thus,

$$\sigma = \sqrt{2}\Phi^{-1}\left(\frac{1+Gini}{2}\right).$$

Solving this equation gives a standard deviation of 0.634 for the US and of 0.39 for Scandinavian countries, or a ratio of 1.63. Assuming that this ratio also applies for the distribution of consumption, one obtains $\triangle \ln \sigma = \ln 1.63$.

Given the estimated values for $\triangle \ln \sigma$ and $\triangle \ln (1+\lambda)$, we can approximate $\varepsilon_2 \simeq \frac{\triangle \ln \sigma}{\triangle \ln(1+\lambda)} = \frac{\ln 1.63}{\ln 1.25} = 2.19$. Again, this elasticity and γ is all that is needed to compute optimal inequality.

Estimates of A and ε_1 are needed to compute optimal choices λ^* and μ^* and welfare costs of current choices, λ^c. To compute A, we assume that the observed levels of σ, μ, and λ, $(\sigma_{y0}, \mu_0, 0)$ lie on the technological frontier. In that case, A can be solved from (18), given ε_1, as $A = \sigma_{y0} (1+\mu_0)^{-\varepsilon_1}$.

We provide two different estimates of ε_1. The first one assumes that the baseline growth rate is optimal. This is likely to provide a robust estimate of ε_1 since long-term growth does not seems to differ much across countries (Klenow and Rodriguez-Clare 2004), suggesting current growth rates are approximately optimal. In that case equation (22) can be used to solve for ε_1 as:

$$\varepsilon_1^a = \frac{\varepsilon_2}{\beta^{-1} (1+\mu_0)^{\gamma-1} - 1}.$$

This way of estimatimg ε_1 allows us to compute λ^* and λ^c, now denoted λ_a^*, and λ_a^c to stress their dependence on ε_1^a, using formulas (23) and (24). By construction, the optimal growth rate is μ_0.

An alternative way to estimate ε_1 is to note that it corresponds an elasticity, i.e. $\varepsilon_1 = \frac{\partial \ln \sigma}{\partial \ln(1+\mu)}$ as derived from (18). One possibility to calibrate this parameter is to turn to the empirical research on inequality and growth. Such estimates appear in this literature and often suggest that the relationship between inequality and growth is negative, but these results have recently been challenged.[17] For example, Alesina and Rodrik (1994) as well as Persson and Tabellini (1994) present evidence that the relationship between inequality and growth is negative but Partridge (1997) disputes the robustness of their findings. Deninger and Squire (1998) argue that data in previous studies of this relationship were of low quality. Using improved measures, they find that asset inequality reduces long-term growth but that it is detrimental to the income growth of the poor rather than the rich. Forbes (2000) argues that these studies suffer from omitted-variable bias and that the relationship is positive. This elasticity could also be computed as the percentage difference in inequality relative to the percentage difference in gross growth rates for two similar countries with initially similar levels of consumption

[17] See Benabou (1996) for a survey.

per-capita (so that $\triangle \ln(1+\lambda) = 0$). Since there still does not seem to be a consensus on the sign or magnitude of ε_1 in the empirical literature, we prefer to use the natural experiment offered by the 1945 separation of East and West Germany. These former countries likely had very similar levels of per-capita consumption and inequality by the end of the Second World War. The radically different social regimes imposed after the war caused a major divergence of per-capita consumption and inequality over a period of 44 years. According to Biewen (2000, Table 2) by 1990 the log variance of income per capita (σ_y^2) was 0.23 in West Germany and 0.1150 in East Germany. This implies a ratio $\frac{\sigma_{ywest}}{\sigma_{yeast}} = 1.41$. Assuming the same ratio for the standard deviation of consumption, then $\triangle \ln \sigma = \ln 1.41$.

Moreover, according to Biewen (2000, Tables 1 and 2), in 1990 mean income in West Germany was 1686.6 and 782 in west Germany, a ratio of 2.15. This ratio is partly confirmed by Burda and Hunt (2001, Table 3) who report a ratio of 2.32. This ratio exaggerates the existing gap right before the fall of the Berlin wall in 1989 because GDP in East Germany dropped dramatically after reunification. According to Burda and Hunt (Table 3), GDP in East Germany fell 15.6% and 22.7% in 1990 and 1991 respectively while GDP in West Germany grew 5.7% and 4.6% during the same period. A better estimate would be $\frac{y_{west1989}}{y_{east1989}} = 2.32\frac{(1-0.156)*(1-0.227)}{1.057*1.046} = 1.37$. Using data from the The Penn World Table Mark 5.6, this ratio was 1.95 in 1970 but only 1.28 in 1988. Overall, these numbers suggest a range for the ratio of per-capita consumptions in 1989 of 1.4 to 1.9. Assuming a ratio of per-capita consumptions of 1 in 1945 and 1.65 in 1989, the ratio of annual growth rates is given by $\frac{1+\mu_{west}}{1+\mu_{east}} = 1.65^{1/44}$. This implies that $\triangle \ln(1+\mu) = \ln \frac{1+\mu_w}{1+\mu_e}$. Thus, our second estimate of ε_1 is $\varepsilon_1^b = \frac{\triangle \ln \sigma}{\triangle \ln(1+\mu)} = 30$.

3.3.1 Results

The first row of Table 5 reports the optimal levels of inequality for different degrees of risk aversion. It ranges from 0.676 for $\gamma = 1$ to 0.15 for $\gamma = 20$. These results suggest that the current degree of inequality in the US is optimal if the coefficient of risk aversion is around 2. However, if the coefficient is larger than 2, as the equity-premium puzzle suggests, then the current level of inequality is excessive. If $\gamma = 5$, for example, then the optimal inequality level is similar to that of Scandinavian countries. Figure 4 illustrates both the optimal choice for the US and the actual choice for inequality and growth. At the observed choice, the production possibility frontier and the social indifference curve for inequality and growth are not tangent, indicating that US consumers would benefit from lower inequality.

Table 5 also reports the calibrated values of ε_1^a. They range from 3.4 for $\gamma = 20$ to 43.5 for $\gamma = 1$. The Table also reports the values of λ^* and λ^c associated to ε_1^a labelled λ_a^* and λ_a^c respectively. By construction the optimal growth rate for this case is $\mu_0 = 2.1\%$ regardless of γ. Consistently with the previous finding, if γ is around 2 then current US choices are approximately optimal and its welfare costs are close to zero. But welfare cost of current choices are significant for other values of γ. If $\gamma = 5$ instead, then reducing inequality to an optimal level of $\sigma^* = 0.3$ would cost a 20% reduction in consumption level but would have a welfare gain of 15%.

Finally, Table 5 reports the values of μ^*, λ^*, and λ^c associated with $\varepsilon_1^b = 0.30$, labelled μ_b^*, λ_b^* and λ_b^c respectively. As γ goes to 1, it becomes optimal to sacrifice all consumption ($\lambda_b^* \to -1$) in order to obtain an infinite growth rate of output. Notice that $\varepsilon_1^a \approx \varepsilon_1^b$ for $\gamma = 2$. This reinforces our previous finding that current US choices for inequality level, growth, and consumption levels, are close to optimal if $\gamma = 2$. However, current US choices have significant welfare costs if $\gamma \neq 2$. For the case of $\gamma = 5$, the optimal inequality level is still a level similar to that of Scandinavian countries, $\sigma^* = 0.3$, but now the tradeoff is between inequality and growth rather than inequality and levels. In fact, for $\gamma = 5$, the optimal policy would require leaving consumption levels around the baseline level, but reducing inequality would imply a fall of the growth rate from 2.1% to 0.45%. Whether this scenario is the relevant one or not, an important consequence of this finding is that the gains from economic growth do not always dwarf the costs of inequality, as Lucas suggested. It is very plausible that the burden of inequality is sufficiently large to merit some sacrifice in terms of economic growth.

4 Final Comments

The main lottery individuals face during their lifetime is their place of birth, and the second, their parents. Children that are identical in all respects will start their life with vast differences in resources and opportunities, depending on which country and what family they are born in. How much consumption level and growth would a new born child be willing to give up in order to avoid these birth lotteries? Lucas' (2004) suggests this child would give up very little. Even if the child is born poor, economic growth would help him or her overcome poverty. In this paper we show that, on the contrary, this child may well be willing to give up all growth to avoid birthplace risk, and a large fraction of growth, if not all, to avoid family risk. The critical elements for our results are time discounting and risk aversion. Both factors downplay the role of growth for welfare while risk aversion enhances the benefits of more equal outcomes. A third key factor is the size of the risk involved at birth, which is staggering.

The contribution of this article is to quantify the social cost of inequality under the most standard assumptions made in macroeconomic theory. Our results suggest that societies could greatly benefit from reducing inequality. They also suggest the existence of a 'big tradeoff' between inequality and efficiency, as described by Okun (1975), and help rationalize why societies may not always find it optimal to adopt growth-enhancing institutions, particularly when inequality is large and those institutions may foster inequality.

Societies commonly face major choices between equality on one hand and efficiency and growth on the other hand. The degree of progressivity of the tax system, for example, reveals the willingness to trade equality for efficiency. Other examples are trade liberalization and labor market reforms, which are often regarded as beneficial for economic efficiency but detrimental in terms of equality. Similarly, the extent of law enforcement, illustrated for example in the efforts to crack down on tax evasion or informal markets, is influenced by distributional concerns at the expense of efficiency and growth. Migration policies are also strongly influenced by this tradeoff. Any correct evaluation of institutional and policy choices made by different societies requires the proper assessment of the welfare implications of these choices, and in particular, a careful consideration of the welfare gains of efficiency and growth against the welfare costs of more inequality. Our results suggest that inequality concerns are of the utmost importance and should be explicitly considered in any aggregate evaluation of institutions.

Public discussions of the costs of inequality and the value of redistributive policies are often framed in political than academic discourse and can lead to mistaken impressions of their effect on social welfare. Glaeser (2005) argues that social attitudes towards inequality in the U.S. and Europe are more the result of ideological language from all sides of the political spectrum than reality. We believe that our work provides an important first step in objectively evaluating the costs of inequality in a well-understood welfare framework.

A caveat of our exercise is that we have not fully specified the micro foundations of the technological restrictions for inequality, consumption growth and levels. We postulate a reduced-form technology, and calibrate it using natural experiments rather than cross-country regressions. Some estimates of the inequality-growth tradeoff appear in the empirical literature on inequality and growth, but there is still no consensus on this relationship. Since most countries seem to share the same long run growth rate, we suspect that the main tradeoff is not between inequality and growth, but inequality and consumption *levels*. In future research, we hope to improve our measurement of these technological constraints using panel data for inequality and consumption levels.

Finally, our exercise suggests that the following is the proper ranking of issues in macroeconomics from the point of view of their potential social welfare impact: cross-country inequality, within-country inequality, economic growth, and business cycles.

Figure 1: Log of Average World Per Capita Consumption, 1960-2000

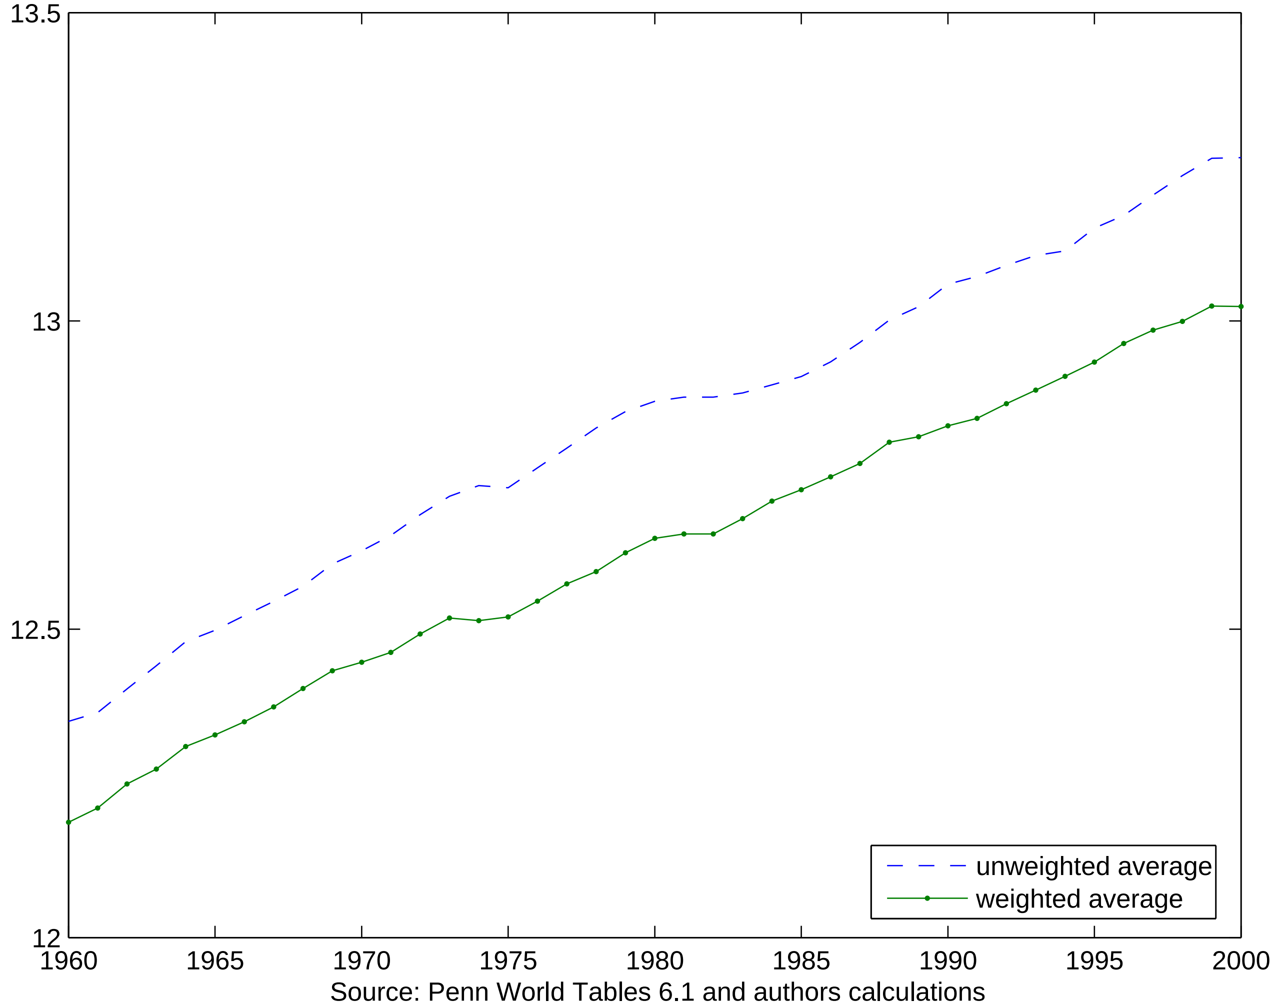

Figure 2: Standard Deviation of Log of Average World Per Capita Consumption, 1960-2000

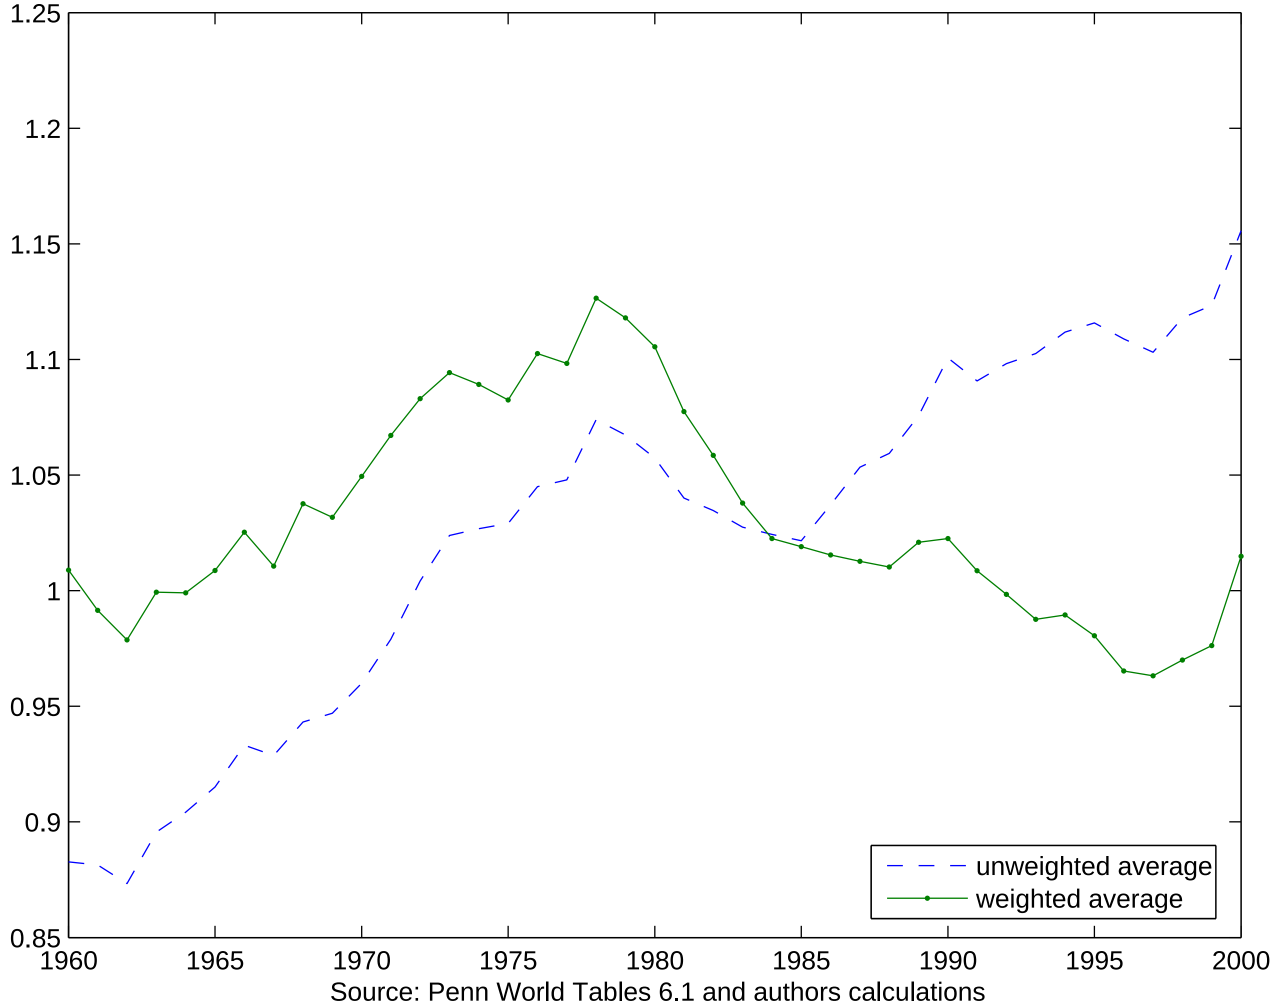

Table 1: Parameter values

$\mu_0 = 2.1\%$ (average per-capita consumption growth rate),
$\sigma_x = 1$ (standard deviation of per-capita consumption across countries),
$\sigma_y = 0.5$ (standard deviation of individual consumption within countries),
$\gamma = [1, 2, 5, 10, 20]$ (coefficient of risk aversion),
$\beta = [0.952, 0.952, 0.949, 0.939, 0.904]$ (discount factor).

Table 2: Welfare Measures - Standard Formulation

γ	1	2	5	10	20
β	0.9521	0.9517	0.9488	0.9394	0.9037
λ_μ	0.5107	0.4050	0.2548	0.1545	0.0766
$\lambda_{1\%}$	0.2155	0.1571	0.0842	0.0449	0.0205
λ_x	-0.3965	-0.6321	-0.9179	-0.9933	-1.000
λ_y	-0.1186	-0.2212	-0.4647	-0.7135	-0.9179
λ_0	-0.4681	-0.7135	-0.9561	-0.9981	-1.000

Note: λ_μ is the welfare gain of economic growth; $\lambda_{1\%}$ is the welfare gain associated with one additional percentage point of growth; λ_x is the welfare cost of cross-country consumption inequality; λ_y is the welfare cost of within-country consumption inequality; λ_0 is the welfare cost of total inequality.

Table 3: Welfare Measures - Alternative Formulation

γ	1.0001	2	5	10	20
β	0.9521	0.9517	0.9488	0.9394	0.9037
$\widehat{\lambda}_{\mu}$	-0.3381	-0.2883	-0.2031	-0.1338	-0.0711
$\widehat{\lambda}_{1\%}$	-0.1773	-0.1358	-0.0777	-0.0430	-0.0201
$\widehat{\lambda}_{x}$	0.6570	1.718	11.18	147.4	22052
$\widehat{\lambda}_{y}$	0.1346	0.2840	0.8682	2.4903	11.182
$\widehat{\lambda}_{0}$	0.8800	2.4903	21.759	517.01	268336

Note: $\hat{\lambda}_{\mu}$ is the welfare cost of zero growth; $\hat{\lambda}_{x}$ is the welfare gain of no cross-country consumption inequality; $\hat{\lambda}_{y}$ is the welfare gain of no within-country consumption inequality; $\hat{\lambda}_{0}$ is the welfare gain of no inequality.

Table 4: Welfare Measures - Growth and Inequality Equivalents

γ	1	2	5	10	20
β	0.9521	0.9517	0.9488	0.9394	0.9037
μ_x	-0.0256	-0.0449	-0.0340	-0.0279	-0.0263
μ_y	-0.0065	-0.0162	-0.0315	-0.0279	-0.0263
μ_0	-0.0319	-0.0504	-0.0340	-0.0279	-0.0263
θ_x	0.8170	0.3400	0.0908	0.0287	0.0074
θ_y	3.2679	1.3602	0.3631	0.1149	0.0295
θ_0	0.6536	0.2720	0.0726	0.0230	0.0059
$\overline{\lambda}_x$	-0.0883	-0.4831	-0.8970	-0.9922	-1.000
$\overline{\lambda}_y$	0.3315	0.0942	-0.3284	-0.6692	-0.9116
$\overline{\lambda}_0$	-0.1964	-0.5975	-0.9449	-0.9978	-1.0000
MRS_1^x	0.0521	0.1488	0.7417	2.8942	13.116
MRS_1^y	0.0260	0.0744	0.3709	1.4471	6.5582
MRS_1^0	0.0582	0.1663	0.8293	3.2358	14.664
MRS_2^x	1.0000	2.0000	5.0000	10.000	20.000
MRS_2^y	0.5000	1.0000	2.5000	5.0000	10.000
MRS_2^0	1.1292	2.2361	5.5902	11.180	22.360

Note: μ_x, μ_y and μ_0 are the points of economic growth that the planner would be willing to accept in exchange for eliminating cross-country inequality, within-country inequality, and total inequality; θ_x, θ_y and θ_0 are the percentage reduction in cross-country inequality, within country inequality, and total inequality respectively that a planner would be willing to accept in exchange for zero economic growth; $\overline{\lambda}_x$, $\overline{\lambda}_y$ and $\overline{\lambda}_0$ is the welfare cost (gain if negative) of eliminating growth and cross-country inequality, growth and within-country inequality, and growth and total inequality respectively; $\text{MRS}_1^i = \frac{\partial\sigma}{\partial\mu}$ and $\text{MRS}_2^i = \frac{\partial\sigma}{\partial\lambda}$ for i={x,y,0}.

Figure 3: Social Indifference Curves

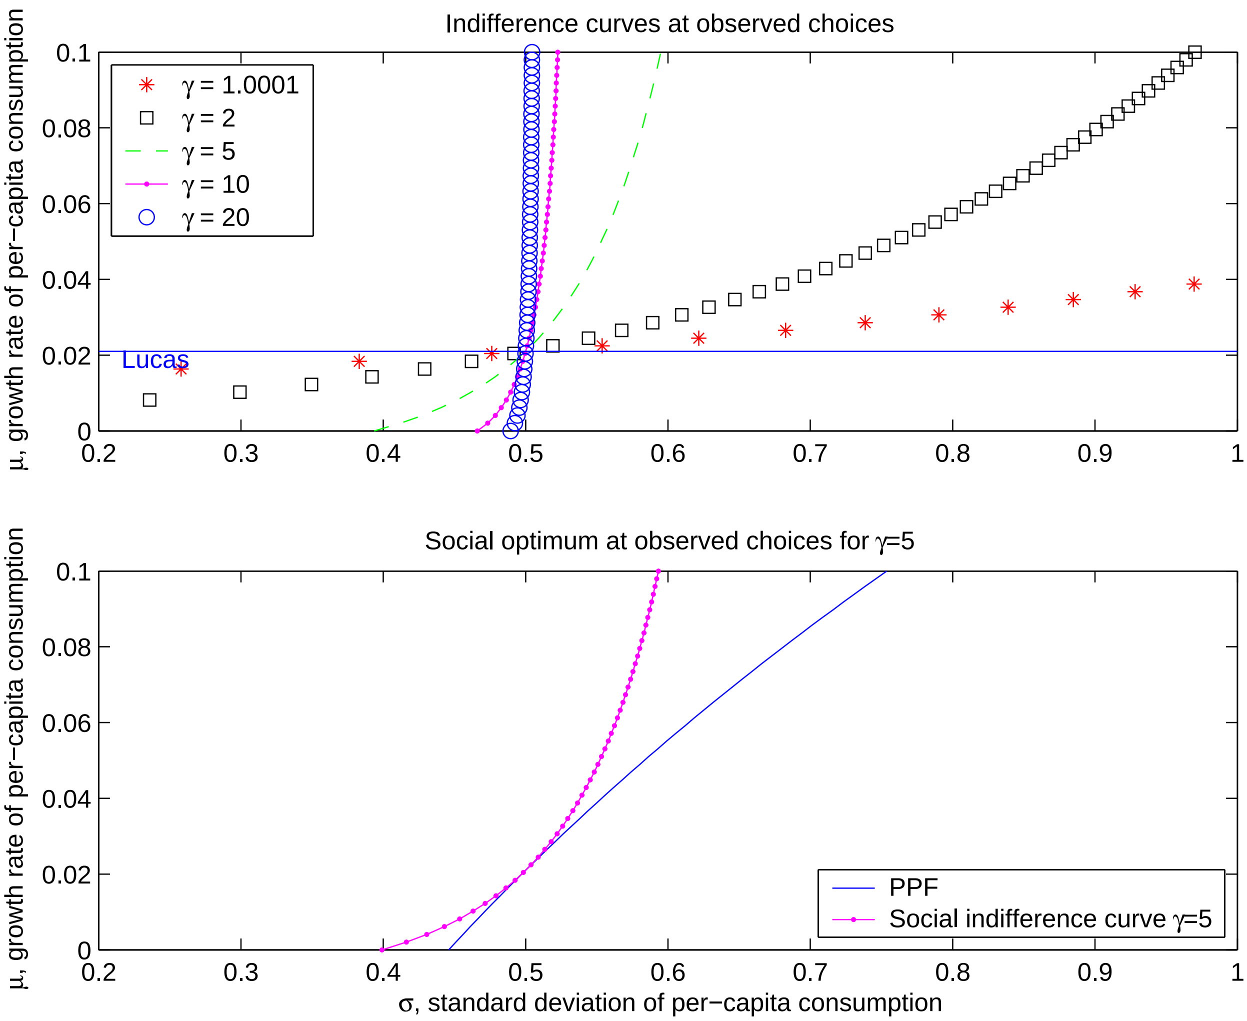

Table 5: Optimal Inequality, Growth, and Consumption Levels

γ	1	2	5	10	20
β	0.9521	0.9517	0.9488	0.9394	0.9037
σ^*	0.6724	0.4778	0.3022	0.2137	0.1511
ε_1^a	43.38	30.06	15.07	7.72	3.41
μ_a^*	0.0210	0.0210	0.0210	0.0210	0.0210
λ_a^*	0.1448	-0.0205	-0.2054	-0.3217	-0.4210
λ_a^c	-0.0327	-0.0010	-0.1536	-0.4693	-0.8219
ε_1^b	30.00	30.00	30.00	30.00	30.00
μ_b^*	∞	0.0211	0.0045	0.0009	-0.0016
λ_b^*	-1.0000	-0.0224	-0.0066	-0.1091	-0.2130
λ_b^c	-1.0000	-0.0010	-0.2089	-0.5394	-0.8563

Note: σ^* is optimal inequality; $\varepsilon_1^a = \frac{\partial \ln \sigma}{\partial \ln(1+\lambda)}$, μ_a^* =optimal growth, λ_a^* =optimal consumption level, and λ_a^c=welfare cost of baseline, values computed assuming $\sigma^* = \sigma_{y0}$; $\varepsilon_1^b = \frac{\partial \ln \sigma}{\partial \ln(1+\lambda)}$, μ_b^* =optimal growth, λ_b^* =optimal consumption level, and λ_b^c=welfare cost of baseline, values computed using German data.

Figure 4: Optimal Inequality and Growth

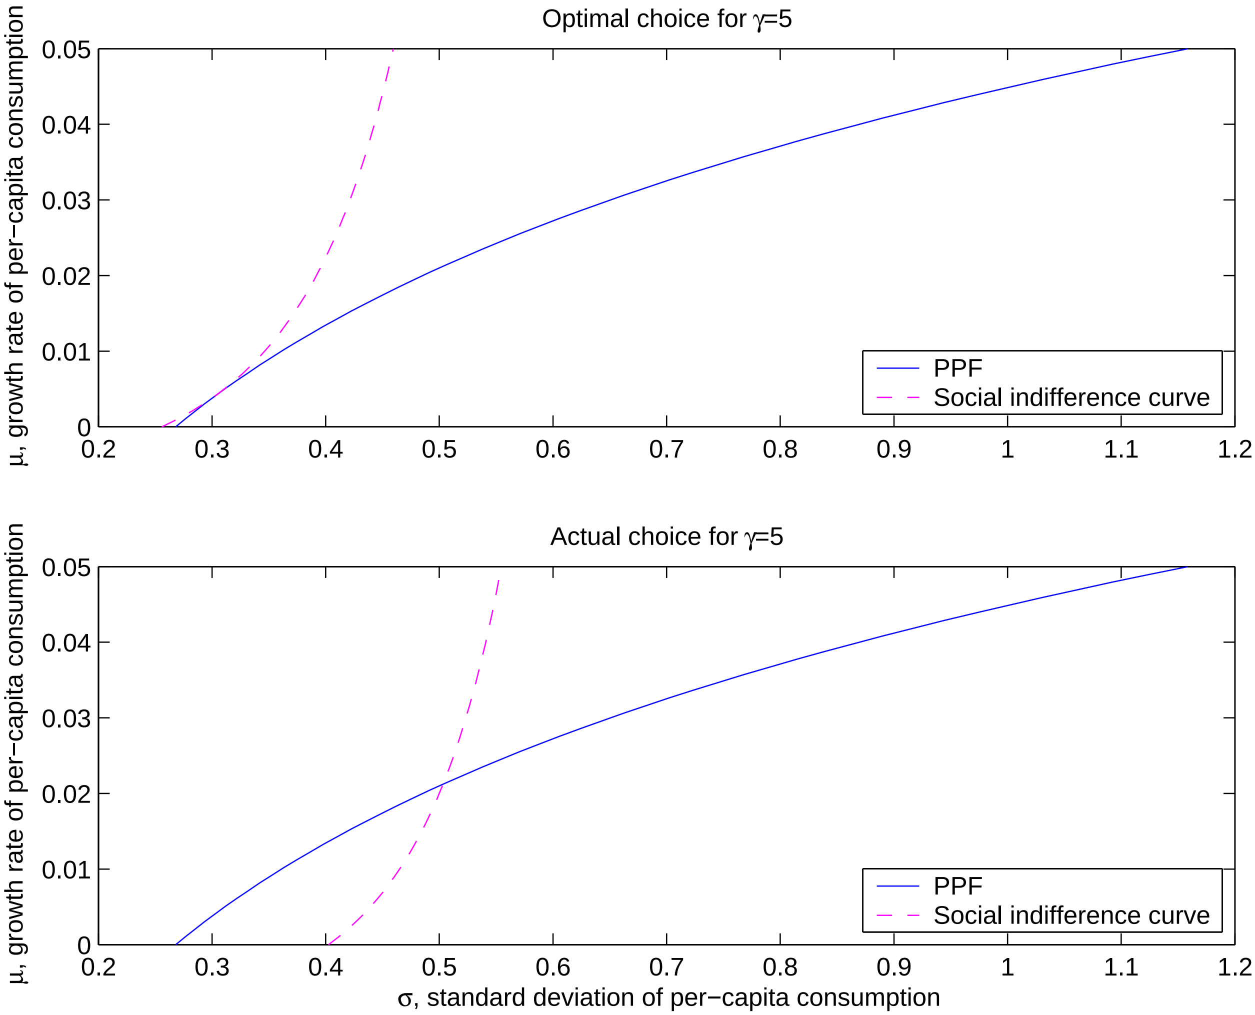

References

Aaberge et al. (2002), "Income Inequality and Income Mobility in the Scandinavian Countries Compared to the United States." Review of Income and Wealth, 48(4).

Acemoglu, Daron (2005), "Modeling Inefficient Institutions." mimeo, MIT.

Aiyagari, S. Rao. (1994) "Uninsured Idiosyncratic Risk and Aggregate Savings." *Quarterly Journal of Economics* 109: 659-684.

Atkeson, Andrew and Lucas, Robert E. (1992) "On efficient Distribution with Private Information," *Review of Economic Studies* 59(3):427-453.

Auerbach, Alan J.; Hassett, Kevin A. (1999) "A New Measure of Horizontal Equity," NBER Working Paper 7035.

Barlevy, G. (2005) "The Costs of Business Cycles and the Benefits of Stabilization" *Economic Perspectives*, 1Q: 1-18.

Beaudry, P. and van Wincoop, Eric. (1996) "The Intertemporal Elasticity of Substitution: An Exploration Using a US Panel of State Data," *Economica*, 63: 495-512.

Beaudry, P. and Pages, C. (2001) "The Costs of Business Cycles and the Stabilization Value of Unemployment Insurance," *European Economic Review*, 45: 1545-1572.

Benabou, R.(1996) "Inequality and Growth," *NBER Macroeconomics Annual*, MA: MIT PRess, 11-74.

Bourguignon, Francois and Christian Morrison (2002) "Inequality among world citizens: 1820-1992," *American Economic Review*, 92(4), September

Campbell, J. and Mankiw, G. (1989) "Consumption, income, and interest rates: reinterpreting the time series evidence." In O. Blanchard and S. Fisher (eds), *NBER Macroeconomic Annual 1989:* 185-216. Cambridge, Mass.: NBER.

Castañeda, Ana; Días-Giménez, Javier; and Ríos-Rull, José-Víctor. ”Accounting for the U.S. Earnings and Wealth Inequality.” Journal of Political Economy 111 (August 2003): 818-857.

Costantiniades, G. M and Duffie, D. (1996) ”Asset Prices with Heterogeneous Consumers,” *Journal of Political Economy* 104(2):219-240.

Deninger, K and Squire, L. (1998) “New Ways of Looking at Old Issues: Inequality and Growth,” *Journal of Development Economics*, 57: 259-287.

Farhi, E. and Werning I. (2005) “Inequality, social discounting, and state taxation,” Mimeo.

Forbes, Kristin J. (2000) “A Reassessment of the Relationship between Inequality and Growth.” *The American Economic Review,* 90(4): 869-887.

Glaeser, Edward L. (2005) “Inequality.” NBER Working Paper 11511.

Gourinchas, Pierre-Olivier & Jonathan A. Parker (2002) ”Consumption Over the Life Cycle,” *Econometrica* 70(1): 47-89.

Hall, R.E. *(1988)* “Intertemporal substitution in consumption,” *Journal of Political Economy*, 96: 339-57.

Harsanyi, John C. (1953) ”Cardinal Utility in Welfare Economics and in the Theory of Risk Taking,” *The Journal of Political Economy* 61(4): 434-5.

Harsanyi, John C. (1955) ”Cardinal Welfare, Individualistic Ethics, and Interpresonal Comparisons of Utility,” *The Journal of Political Economy* 63(4): 309-321.

Harsanyi, John C. (1975) ”Nonlinear Social Welfare Functions: Do Welfare Economists Have a Special Exemption From Bayesian Rationality?” *Theory and Decision* 6: 311-32.

Harsanyi, John C. (1982) ”Morality and the theory of rational behavior,” in *Utilitarianism and Beyond*, Edited by Amartya Sen and Bernard Williams, Cambridge University Press, London, 1982.

Kaplow, Louis (2003) “Concavity of Utility, Concavity of Welfare, and Redistribution of Income,” NBER Working Paper 10005

King, R. and S. Rebelo (2000) "Resuscitating Real Business Cycles", Chapter 14, Volume 1B, Handbook of Macroeconomics, J. Taylor and M. Woodford eds, North Holland.

Klenow, Peter J. and Rodríguez-Clare, Andrés. "Externalities and Growth." *NBER Working paper* 11009, December 2004.

Kocherlakota, Narayana R. (1996) "The Equity Premium; It's Still a Puzzle," *Journal of Economic Literature*, 34(1): 42-71.

Krueger, Dirk; and Perri, Fabrizio (2002) "Does Income Inequality Lead to Consumption Inequality? Evidence and Theory," NBER Working Paper 9202.

Krusell, P., and Smith, A. (1999) "On the Welfare Effects of Eliminating Business CYcles," *Review of Economic Dynamics*, 2(1): 245-272.

Kutznets, Simon (1955) "Economic Growth and Income Inequality," *American Economic Review*, 45(1): 1-28.

Lucas, Robert E, (1988) "On the Mechanics of Economic Development" *Journal of Monetary Economics 22:* 3-42.

Lucas, Robert E, *Lectures on Economic Growth*, Harvard University Press, Cambridge, 2002.

Lucas, Robert E. (2004) "The Industrial Revolution: Past and Future," *The Region* (May), Federal Reserve Bank of Minneapolist, pages 5-20.

Lucas, Robert E. (1987) *Model of Business Cycles*, Basil Blackwell, New York, 1987.

Mirrleess, J.A. (1982) "The economics uses of utilitarianism" in *Utilitarianism and Beyond*, Edited by Amartya Sen and Bernard Williams, Cambridge University Press, London, 1982.

Okun, Arthur M. (1975) *Equality and Efficiency: The Big Tradeoff,* The Brookings Institutions, Washington.

Partridge, Mark D. (1997) "Is Inequality Harmful for Growth? Comment." *The American Economic Review,* 87(5): 1019-1032.

Pattanaik, Prasanta K. (1968) "Risk, Impersonality, and the Social Welfare Function," *Journal of Political Economy*, 76(6):1152-1169.

Persson, Torsten; and Tabellini, Guido. (1994) "Is Inequality Harmful for Growth? " *The American Economic Review,* 84(3): 600-621.

Phelan, C. (2003) "Opportunity and social mobility" *Staff Report,* FRB Minneapolis.

Ng, Y.-K. (1975) "Bentham or Bergson? Finite Sensiblity, Utility Functions, and Social Welfare Functions," *Review of Economic Studies* 42: 545-69.

Ricardo, David (1951) Works and Correspondence. Edited by Pierro Sraffa. 10 vols. Cambridge: Cambridge University Press.

Sala-I-Martin (2002) "The Disturbing "Rise" of World Income Inequality," NBER Working paper 8904, April 2002.

Sen, A. K. (1970) *Collective Choice and Social Welfare*, Hoden Day, San Francisco.

Sleet, Christopher; and Yeltekin, Sevin (2005) "Social credibility, social patience and long run inequality" mimeo, July.

Storesletten, Kjetil; Telmer, Chris I.; and Yaron, Amir (2004) "Cyclical Dynamics in Idiosyncratic Labor-Market Risk." Journal of Political Economy 112(3):695-717.

Tallarini, T. (2000) "Risk-Sensitive Real Business Cycles," *Journal of Monetary Economics,* 45: 507-532.

Vickrey, William (1960) "Utility, Strategy, and Social Decision Rules," *Quarterly Journal of Economics*, 74 (4):507-535.